



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 1 107 123 A1

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
13.06.2001 Bulletin 2001/24

(51) Int Cl.7: G06F 12/08

(21) Application number: 99403042.7

(22) Date of filing: 06.12.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Chauvel, Gérard
06600 Antibes (FR)
• Lasserre, Serge
83600 Frejus (FR)
• d'Inverno, Dominique Benoit Jacques
06270 Villeneuve-Loubet (FR)

(71) Applicants:
• TEXAS INSTRUMENTS FRANCE
06271 Villeneuve Loubet Cédex (FR)
Designated Contracting States:
FR
• Texas Instruments Incorporated
Dallas, Texas 75251 (US)
Designated Contracting States:
BE CH DE DK ES FI GB GR IE IT LI LU MC PT AT
CY

(74) Representative: Obolensky, Michel et al
c/o CABINET LAVOIX
2, place d'Estienne d'Orves
75441 Paris Cédex 09 (FR)

(54) Smart cache

(57) A cache architecture (16) for use in a processing device includes a RAM set cache for caching a contiguous block of main memory. The RAM set cache can be used in conjunction with other cache types, such as a set associative cache or a direct mapped cache. A register (32) defines a starting address for the contiguous block of main memory. The data array (38) associated with the RAM set may be filled on a line-by-line basis, as lines are requested by the processing core, or on a

set-fill basis which fills the data array (38) when the starting address is loaded into the register (32). As addresses are received from the processing core, hit/miss logic (46) the starting address register (32), a global valid bit (34), line valid bits (37) and control bits (24, 26) are used to determine whether the data is present in the RAM set or whether the data must be loaded from main memory. The hit/miss logic (46) also determines whether a line should be loaded into the RAM set data array (38) or in the associated cache.




Fig. 2

Description

CROSS-REFERENCE TO RELATED APPLICATIONS

5 [0001] Not Applicable

STATEMENT OF FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

10 [0002] Not Applicable

BACKGROUND OF THE INVENTION

1. TECHNICAL FIELD

15 [0003] This invention relates in general to processing devices and, more particularly, to a cache architecture for a processing device.

2. DESCRIPTION OF THE RELATED ART

20 [0004] Most processing devices use a cache architecture to increase the speed of retrieving information from a main memory. A cache memory is a high speed memory that is situated between the processing core of a processing device and the main memory. The main memory is generally much larger than the cache, but also significantly slower. Each time the processing core requests information from the main memory, the cache controller checks the cache memory to determine whether the address being accessed is currently in the cache memory. If so, the information is retrieved from the faster cache memory instead of the slower main memory. If the information is not in the cache, the main memory is accessed, and the cache memory is updated with the information.

25 [0005] As processing cores increase in speed relative to memory designs, the efficiency of the cache architecture becomes more significant. One way to increase efficiency is to increase the size of the cache. Since a larger cache memory can store more information, the likelihood of a cache hit is similarly increased. In most cases, however, increasing cache size has diminishing returns after a certain point. Further, increasing the cache size will increase the size of the chip (assuming the cache is integrated with the processing core). Even more importantly, access time will be increased, defeating the initial purpose of the cache. Accordingly, merely increasing the size of a cache will in many cases not produce worthwhile results.

30 [0006] In many devices, certain routines will have critical time constraints or will otherwise need a predictable execution time. In these cases, it can be critical to eliminate latencies due to cache misses. Some cache systems provide mechanisms for locking entries in a cache, so that the cache entries will not be overwritten as other locations are accessed. This mechanism is useful for entries that will be used repeatedly; however, locking entries of a cache reduces the size and associativity of the cache. For instance, in a 2-way set associative cache, locking some entries will result in a portion of the cache acting as a direct map, greatly reducing the efficiency of the cache. A similar solution uses a local memory working in parallel with the cache system. This solution requires address decoding for the local memory and a cache disabling mechanism, which can result in latencies. Further, while an implementation with a local RAM may work with routines specifically written to use the local RAM, other routines, specifically OS (operating system) routines not written in anticipation of the specific local RAM configuration will not be able to control the local RAM in the manner that the cache is controlled.

45 [0007] Therefore, a need has arisen for a cache architecture that increases cache performance and predictability.

BRIEF SUMMARY OF THE INVENTION

50 [0008] In a first embodiment of the present invention, a processing device comprising a processing core having circuitry for generating addresses to access a main memory and an n -way cache. The n -way cache comprises n data memories each having a plurality of entries for storing information from the main memory, one or more tag memories for storing address information identifying a main memory address associated with each of the entries in a corresponding data memory, a plurality of tag registers for storing address information defining a contiguous block of main memory addresses, each tag register associated with a corresponding data memory, and control circuitry for defining a cache association between each data memory and either a tag memory or a tag register and selectively accessing each data memory in response to an address from the processing core based on the cache association.

55 [0009] In a second embodiment of the present invention, a processing device comprises a processing core having circuitry for generating addresses to access a main memory, a first, n -way, cache subsystem, where n is greater than

or equal to 1, and a second, m -way, cache system, where m is greater two or equal to 1. The first cache system comprises n data memories each having a plurality of entries for storing information from the main memory and n tag memories for storing address information identifying a main memory address associated with each of the entries in a corresponding one of the n data memories. The second cache subsystem comprises m data memories each having a plurality of entries for storing information from the main memory and m tag registers, each storing address information defining a contiguous block of main memory addresses mapped to a corresponding one of the m data memories. Logic determines cache hits in the first and second cache subsystems, where hits from the second cache subsystem have precedence over hits from the first subsystem.

[0010] The present invention provides significant advantages over the prior art. First, the RAM set cache (mapped to a contiguous block of main memory addresses) can significantly improve the operation of a processing device performing real-time operations, since a desired block of code can be stored in the RAM set cache for fast retrieval. Second, there is no extra penalty for accessing a larger data memory for a RAM set cache, as long as the access time of the RAM set is not bigger than the access time of the standard cache. Third, the addition of one or more RAM set caches can be provided with a minimal amount of circuitry over a conventional cache. Fourth, the RAM set caches can be configured in a very flexible manner with other caches, such as a set associative or direct map cache, as desired. Fifth, the RAM set cache provides advantages over a local RAM, because a separate mechanism for loading the data memory is not necessary for the RAM set cache and no specific address decoding in serial with the memory access time is required. Sixth, the cache can be controlled by the OS or other software in the same manner as an ordinary cache - loading, flushing, line invalidation, and so on, can be performed by the software without knowledge of the specific architecture of the cache, or with minor modifications to a driver for the OS.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0011] For a more complete understanding of the present invention, and the advantages thereof, reference is now made to the following descriptions taken in conjunction with the accompanying drawings, in which:

Figure 1 illustrates a block diagram of a processing device incorporating a cache;

Figure 2 illustrates a block diagram of a preferred embodiment of a cache architecture;

Figure 3 is a diagram showing the mapping of a portion of main memory onto a RAM set cache; and

Figure 4 illustrates a flow diagram describing operation of the hit/ miss logic of Figure 2.

DETAILED DESCRIPTION OF THE INVENTION

[0012] The present invention is best understood in relation to Figures 1-4 of the drawings, like numerals being used for like elements of the various drawings.

[0013] Figure 1 illustrates a block diagram of a processing device 10. Processing device 10 includes a processing core 12, data memory 14, instruction cache 16, and subsystem memory interface 18. Subsystem memory interface 18 interfaces with main memory 20, which is typically an external memory.

[0014] As described in greater detail below, in the preferred embodiment, the instruction cache is a 3-way cache with one cache way being a "RAM set" cache memory. The RAM set cache is designed to cache a contiguous block of memory starting from a chosen main memory address location. The other two cache ways can be configured as RAM set cache memories, or use another architecture. For example, the instruction cache 16 could be configured as one RAM set cache and a 2-way set associative cache.

[0015] In operation, the processing core 12 accesses main memory 20 within a given address space. If the information at a requested address in main memory is also stored in the instruction cache 16, the data is retrieved from the instruction cache. If information for the requested address is not stored in the instruction cache, the information is retrieved from the main memory 20 and the instruction cache is updated with the retrieved information.

[0016] Figure 2 illustrates a more detailed block diagram of the instruction cache 16, in an embodiment with a RAM set cache and a two-way set associative cache.

[0017] A cache controller 22 controls operation of the instruction cache 16. Cache controller 22 includes four status bits: RAM_fill_mode 24, Cache_Enable 26, DM/2SA 28 and Full_RAM_base 30. Cache controller 22 is coupled to Full_Set_Tag registers 32 (individually referenced as registers 32₁ through 32₃), Global_Valid bits 34 (individually referenced as bits 34₁ through 34₃), tag memories 36 (individually referenced as tag memories 36₂ and 36₃), valid entry bit arrays 37 (individually referenced as bit arrays 37₁ through 37₃) and data arrays 38 (individually referenced as data arrays 38₁ through 38₃). Comparators 40 (individually referenced as comparators 40₁ through 40₃) are coupled to

EP 1 107 123 A1

respective Full_Set_Tag registers 32. Comparators 42 (individually referenced as comparators 42₂ and 42₃) are coupled to respective tag memories 36. Output buffers 44 (individually referenced as buffers 44₁ through 44₃) are coupled to respective data arrays 38. Hit/Miss logic 46 (individually referenced as logic 46₁ through 46₃) is coupled to comparators 40, global valid bits 34, valid bits 37, RAM_fill_mode bit 24 and Cache_Enable bit 26.

[0018] In operation, instruction cache 16 is configured using the control bits 24, 26, 28 and 30. The Cache_Enable 26 allows the instruction cache to be enabled or disabled, as in standard cache architecture. If the instruction cache 16 is disabled (Cache_Enable=0), instruction read accesses are performed on the main memory 20 via the subsystem memory interface 18, without using the instruction cache 16. If the instruction cache is enabled (Cache_Enable=1), instructions are executed from the instruction cache 16, in cases where such instructions are present in the instruction cache 16. If a miss occurs, a line (for example, 16 bytes) is fetched from main memory 20 and provided to the core 12. This is also standard cache behavior.

[0019] The size of the data array 38₁ can be different than the size of the data arrays 38_{2,3} for the other ways of the cache. For illustration purposes, it will be assumed that data arrays 38₂ and 38₃ are each 8 Kbytes in size, configured as 512 lines, with each line holding eight two-byte instructions. Data array 38₁ is 16 Kbytes in size, configured as 1024 lines, each line holding eight two byte instructions. ADDR[L] is used to address one line of the data array 38 and valid bit array 37 (and tag memory 36, where applicable). Accordingly, for the 1024-line first way, ADDR[L] will include bits [13:4] of an address from the core. For the 512-line second and third ways, ADDR[L] will include bits [12:4] of an address from the core. ADDR[H] defines which set is mapped to a line. Hence, assuming a 4 Gbyte (2 Gword) address space, ADDR[H] uses bits [31:14] of an address from the core for the first way and uses bits [31:13] for each of the second and third ways of the cache 16.

[0020] The tag memories 36 and comparators 42 are used for a 2-way set associative cache. When the core 12 performs a memory access, the tag memories 36 are accessed at the low order bits of the address (ADDR[L]). The tag memory locations store the high order address bits of the main memory address of the information stored in a corresponding line of the data array 38. These high order address bits are compared with the high order address bits (ADDR[H]) of the address from the core 12. If the ADDR[H] matches the contents of the tag memory at ADDR[L], a hit occurs if the valid bit associated with the low order bits (V(ADDR[L])) indicates that the cache entry is valid. If there is a cache hit, the data from the corresponding data array 38 at ADDR[L] may be provided to the core 12 by enabling the proper output buffer 44. As described below, data from the 2-way cache is presented to the core 12 only if there is a miss in the RAM set cache. By itself, the operation of the 2-way set associative cache and the direct map cache can be conventional and is not affected by the RAM set cache. Other cache techniques could also be used in conjunction with the RAM set cache.

[0021] The RAM set cache stores a contiguous block of main memory 20 starting at an address defined by the Full_set_tag register 32 for the RAM set. This block of information is mapped to the corresponding data array 38 of the RAM set. Only the high order bits of the starting address are stored in the Full_set_tag register 32. Figure 3 illustrates this mapping for a single RAM set. As shown, the contents of Full_set_tag register 32₁ defines the starting address for a contiguous block of memory cached in data array 38₁.

[0022] A RAM set miss occurs when the high order bits of the address from the core 12 do not match the contents of the Full_set_TAG register 32 or the global valid bit equals "0". In either case, when there is a RAM set miss, the instruction cache 16 behaves like a normal 2-way cache logic - if there is a hit in the 2-way cache, then an instruction is presented to the core 12 from the 2-way set associative cache; otherwise the instruction is retrieved from main memory 20.

[0023] A RAM set hit situation occurs when the high order bits of the address from the core 12 match the contents of the Full_set_TAG register 32 and the global valid bit equals "1" (the setting of the global valid bit is described in greater detail hereinbelow). The RAM set comparison has the highest priority by default. A hit situation indicates that the requested instruction is mapped to the RAM set. If the Valid entry bit 37 corresponding to the line containing the instruction is set to "1", the logic 40 generates a hit-hit signal, because the address hit the RAM set and the instruction is present in the RAM set. If the corresponding valid bit of the RAM set entry 37 is "0", the logic generates a hit-miss because the address hit the RAM set but the instruction is not yet present in the RAM set. In that case, the instruction is fetched from main memory 20 and is loaded into the data array 38 of the RAM set. A hit in the RAM set takes precedence over the normal cache logic. The standard logic of the 2-way cache always generates a miss when the RAM set logic generates a hit. Information can reside in both the RAM set and the 2-way cache without causing any misbehavior; the duplicated cache entry in the 2-way cache will eventually be evicted by the replacement mechanism of the two-way cache, because it will never be used.

[0024] To set up a RAM set for operation, the Full_set_tag register 32 must be loaded with the start address (set_start_addr) and the RAM_fill_mode bit 24 must be configured to a desired fill mode. The circuitry for filling the cache can be the same as that used to fill lines of the set associative cache. In the preferred embodiment, two fill modes are provided: a line-by-line fill mode or set fill mode.

[0025] For a line-by-line fill (RAM_fill_mode=0), the global valid bit 34 is set to "1" and the valid entry bits 37 are set

to "0" when the Full_set_tag register 32 is loaded with the starting address. At this point, the data array 38 is empty (it is assumed that the Cache_Enable bit 26 is set to "1" to allow operation of the instruction cache). Upon receiving an address from the core 12, a valid entry bit 37 is selected based on the low order bits of the address. As provide above, if the RAM set is 16Kbytes in size, organized as an array of 1K x 16 bytes, where 16 bytes is equivalent to a block line in the associated 2 way cache, the Full_set_TAG register will store 18 bits [31:14] of the starting address (set_start_addr). The address indexing each entry of the RAM set (ADDR[L]) will have 10 bits [13:4] and the instruction address, used to access one instruction in the line, will have 3 bits [3:1] (assuming instructions are 2 bytes wide). A line of the data array 38 (at ADDR[L]) is loaded from main memory 20 each time that a hit-miss situation occurs because (1) the comparator 40 determines a match between ADDR[H] and Set_start_addr, (2) the Global valid bit 34 is set to "1" and (3) the valid bit 37 associated with the line at ADDR[L] is cleared (V[ADDR[L]]="0"). This situation indicates that the selected line is mapped to the RAM set, but has not yet been loaded into the RAM set's data array 38. When the line is loaded into the data array 38 from main memory 20, the valid bit 37 corresponding to the line is set to "1". This loading procedure has the same time penalty as a normal cache line load, but the entry will stay in the RAM set like a locked entry and, therefore, the processing device will not be penalized on a subsequent access.

[0026] On the other hand, if a set fill (RAM_fill_mode) is chosen, the global valid bit 34 is initially set to "0" and remains "0" after the Full_set_tag register is loaded with the starting address. The valid bits 37 are also set to "0". When the starting address is written to the Full_set_tag register 32, the associated data array 38 is filled through a DMA (direct memory access) process. As each line is loaded from main memory 20, the valid entry bit 37 corresponding to the line is set to "1". After the data array 38 has been completely loaded from main memory 20, the global valid bit 34 is set to "1". This initialization of the data array 38 takes longer than the line-by-line fill mode, but all critical real-time routines are available after initialization and the system latency is deterministic. After the RAM set is initialized in set fill mode, there will never be a miss on code mapped to the RAM set, even on the first access.

[0027] In either set-fill or line-by-line fill modes, the contents of a RAM set can be changed simply by writing a new Set_start_addr into the Full_set_tag register 32. Writing to this register flushes the contents of the respective set and initiates a load process according to the fill mode. The control circuitry 22 can use the same circuitry for flushing lines of the RAM set cache as is used for the set associative cache. Flushing an entire RAM set cache can be accomplished simply by writing to the appropriate Full_set_tag register 32. Similarly, the control circuitry 22 can use the same circuitry for filling lines of the RAM set cache as is used for the set associative cache. The operation for filling an entire cache in set fill mode is slightly different because multiple lines, rather than a single line, are accessed from the main memory 20 and stored in the data array 38. The RAM set cache can be used with an OS that is not specifically designed to operate with a RAM set cache through the use of minor driver modifications.

[0028] The operation of the Hit/Miss logic is described in connection with the flow chart of Figure 4. In step 50, an address is received from the core 12 in connection with a read operation. If the instruction cache is disabled in step 52, the instruction is retrieved from main memory 20 in step 54. If the cache is enabled, then if either the high order bits of the address from the core (ADDR[H]) do not match the high order bits of the starting address (Set_start_addr) or the global valid bit 34 is set to "0" (step 56), then there is a RAM set miss. In this case, if there is a cache hit in the 2-way set associative cache (step 58), then the information retrieved from the 2-way set associative cache is presented to the core 12. If there is a miss in the 2-way set associative cache, the line is loaded into the 2-way cache.

[0029] Returning again to step 56, if both the high order bits of the address from the core (ADDR[H]) match the high order bits of the starting address (Set_start_addr) and the global valid bit 34 is set to "1", then there is a RAM set hit at the line corresponding to ADDR[L], and the valid entry bits are used to determine whether it is a hit-hit situation (where the requested instruction is present in the RAM set and can be presented to the core 12) or a hit-miss situation (where the requested instruction is mapped to the RAM set, but the information needs to be loaded into the RAM set's data array 38 from the main memory 20). If, in step 64, the valid entry bit 37 for the line indicates that the line is valid (V[ADDR[L]]=1), the instruction is present in the RAM set and is presented to the core 12 through the RAM set's output buffer 44. If, on the other hand, the valid entry bit 37 for the line indicates that the line is not valid (V[ADDR[L]]=0), the line is loaded into the data array 38 of the RAM set from main memory in step 68.

[0030] The flow chart of Figure 4 can be easily implemented using combinational logic.

[0031] In the preferred embodiment, the cache 16 provides flexibility in providing one, two, or three RAM sets, since different applications for the processing device 10 have different real-time requirements. In this embodiment, control bits DM/2SA 28 and Full_RAM_base 30 define the allocation of RAM sets in the 3-way cache architecture. Table I describes the possibilities for the illustrated embodiment.

Table I

Cache Configurations		
Full_RAM_base	DM/2S A	Configuration
0	0	One 2-way set associative cache and one RAM set cache
1	0	One direct map cache and a two set RAM cache
1	1	Three set RAM cache

[0032] The Full_set_tag register 32 uses a number of bits equal to the length of ADDR[H] for the associated way of the cache. Hence, Full_set_tag register 32₁ stores bits [31:14] and Full_set_registers 32₂ and 32₃ store bits [31:13], for the specific data array sizes and configurations defined herein.

[0033] In a 2-way set associative cache, both tag memories 36 are used; in a direct map cache a single tag memory 36 is used. The tag memories 36 are not used for any of the caches configured as RAM set caches. Thus, for the configuration using a single RAM set cache and a 2-way associative cache, the RAM set cache uses Full_set_tag register 32₁, global valid bit 34₁, valid entry bits 37₁, data array 38₁, comparator 40₁, Hit/miss logic 46₁, and output buffer 44₁. The 2-way set associative cache would use tag memories 36₂ and 36₃, valid bits 37₂ and 37₃, data arrays 38₂ and 38₃, Hit/miss logic 46₂ and 46₃, and output buffers 44₂ and 44₃. For a configuration using two RAM sets and a direct mapped cache, the RAM sets would use Full_set_tag registers 32₁ and 32₂, global valid bits 34₁ and 34₂, valid entry bits 37₁ and 37₂, data arrays 38₁ and 38₂, comparators 40₁ and 40₂, Hit/miss logic 46₁ and 46₂, and output buffers 44₁ and 44₂. The direct mapped cache would use tag memory 36₃, valid bits 37₃, data array 38₃, Hit/miss logic 46₃ and output buffer 44₃. For a configuration using three RAM sets, the RAM sets would use Full_set_tag registers 32₁, 32₂ and 32₃, global valid bits 34₁, 34₂ and 34₃, valid entry bits 37₁, 37₂ and 37₃, data arrays 38₁, 38₂ and 38₃, comparators 40₁, 40₂ and 40₃, Hit/miss logic 46₁, 46₂ and 46₃, and output buffers 44₁, 44₂ and 44₃.

[0034] While the embodiment shown provides a 3-way cache, any number of cache-ways could be provided. For example, a 4-way cache could be configured to use and combination of RAM set and set associative, or other, cache architectures. The only additional hardware needed for the additional RAM set cache would be the additional Full_set_tag register and the global valid bit.

[0035] The RAM set cache is compatible with self-modifying code. If the processing core 12 changes an instruction dynamically, the line containing the modified location is flushed from the cache (i.e., its corresponding valid bit 37 is set to "0") in parallel with the write operation to main memory. The next time that the instruction is requested by the core 12, the corresponding valid bit will be set to "0", causing a hit-miss condition. The line containing the requested instruction will be loaded into the RAM set cache from main memory. In the preferred embodiment, instructions are not modified in the cache directly, eliminating the need to update the main memory when replacing a line in the cache.

[0036] While the invention has been discussed in connection with an instruction cache, it could also be used as a data cache, or as a unified instruction/data cache.

[0037] The present invention provides significant advantages over the prior art. First, the RAM set cache can significantly improve the operation of a processing device performing real-time operations, since a desired block of code can be stored in the Ram set for fast retrieval. Second, there is no extra penalty for accessing a larger data array 38 for a RAM set cache, as long as the access time of the RAM set is not bigger than the access time of the 2-way cache. Third, the addition of one or more RAM set caches can be provided with a minimal amount of circuitry over a conventional cache. The only additional circuitry required is one or more Full_Set_Tag registers 32 and the associated global valid bits 34 (the valid bits 37 are necessary for the RAM set(s) only when a line-by-line fill mode is available or self-modifying code is allowed; if the RAM set only supports set fill without self-modifying code, the valid bits 37 would be unnecessary). The valid bits 37 of a set associative or other cache can be used for the RAM set as well, if the valid bits are not included in the tag memory itself. Fourth, the RAM sets can be configured in a very flexible manner with other caches, such as a set associative or direct map cache, as desired. Fifth, the RAM set cache provides advantages over a local RAM, because a separate mechanism for loading the RAM is not necessary for the RAM set cache and no specific address decoding in serial with the memory access time is required. Sixth, the cache can be controlled by the OS or other software in the same manner as an ordinary cache - loading, flushing, line invalidation, and so on, can be performed by the software with minor software adaptation using conventional cache management.

[0038] While the present invention has been shown for a specific embodiment herein, it could be used in a number of implementations. First, the RAM set cache architecture could be used in any type of processing device, including microprocessors, DSPs, mixed analog/digital processors, and co-processors. Second, the sizes of the data arrays could be varied as needed for a certain implementation with minor modifications. For example, it may be desirable to have a RAM set with a larger data array than the set associative cache, or vice-versa, depending upon the size of the

applications with real-time constraints. Third, while the preferred embodiment allows the cache types to be mixed in a flexible manner, it may be preferable in some circumstances to have set cache types, such as a single RAM set and a set-associative cache. Fourth, the architecture used to implement non-RAM set caches (i.e., the set associative and direct mapped caches) could use different architecture from that shown herein; for example, a CAM (content addressable memory) could be used for the Tag memories 36.

[0039] Although the Detailed Description of the invention has been directed to certain exemplary embodiments, various modifications of these embodiments, as well as alternative embodiments, will be suggested to those skilled in the art. The invention encompasses any modifications or alternative embodiments that fall within the scope of the Claims.

Claims

1. A processing device comprising:

a processing core having circuitry for generating addresses to access a main memory;
an n -way cache comprising:

n data memories each having a plurality of entries for storing information from said main memory;
one or more tag memories for storing address information identifying a main memory address associated with each of said entries in a corresponding data memory;
a plurality of tag register (32) for storing address information defining a contiguous block of main memory addresses, each tag register associated with a corresponding data memory; and
control circuitry for defining a cache association between each data memory and either a tag memory or a tag register and selectively accessing each data memories in response to an address from said processing core based on said cache association.

2. The processing device of claim 1 and further comprising a global valid bit associated with each of said tag registers.

3. The processing device of claim 2 and further comprising a valid entry array comprising a valid entry bit corresponding to each entry of a corresponding data memory.

4. The processing device of claim 1 wherein said control circuitry comprises configuration bits for defining the cache association

5. The processing device of claim 1 wherein said control circuitry can define an association between m tag memories and m data memories to form an m -way set associative cache and $n-m$ caches for storing information mapped to respective blocks of main memory addresses.

6. The processing device of claim 1 and further comprising logic to determine the occurrence of a cache hit in said cache.

7. An n -way cache system comprising:

n data memories each having a plurality of entries for storing information from a main memory;
one or more tag memories for storing address information identifying a main memory address associated with each of said entries in a corresponding data memory;
a plurality of tag registers for storing address information defining a contiguous block of main memory addresses, each tag register associated with a corresponding data memory; and
control circuitry for defining a cache association between each data memory and either a tag memory or a tag register and selectively accessing each data memories in response to an address from a processing core based on said cache association.

8. The cache system of claim 7 and further comprising a global valid bit associated with each of said tag registers.

9. The cache system of claim 8 and further comprising a valid entry array comprising a valid entry bit corresponding to each entry of a corresponding data memory.

EP 1 107 123 A1

10. The cache system of claim 7 wherein said control circuitry comprises configuration bits for defining the cache association

11. The cache system of claim 7 wherein said control circuitry can define an association between m tag memories and m data memories to form an m -way set associative cache and n - m caches for storing information mapped to respective blocks of main memory addresses.

12. The cache system of claim 7 and further comprising logic to determine the occurrence of a cache hit in said cache.

13. A processing device comprising:

a processing core having circuitry for generating addresses to access a main memory;
a first n -way cache subsystem, where n is greater than or equal to 1, comprising:

n data memories each having a plurality of entries for storing information from said main memory;
 n tag memories for storing address information identifying a main memory address associated with each of said entries in a corresponding one of said n data memories;

a second m -way cache subsystem, where m is greater than or equal to 1, comprising:

m data memories each having a plurality of entries for storing information from said main memory;
 m tag registers, each storing address information defining a contiguous block of main memory addresses mapped to a corresponding one of said m data memories; and

logic for determining cache hits in said first and second cache subsystems, where hits from said second cache subsystem have precedence over hits from said first subsystem.

14. The processing device of claim 13 and further comprising cache control circuitry for filling one or more lines of one of said cache subsystems after a cache miss in said first and second cache subsystems.

15. The processing device of claim 13 wherein said first cache subsystem comprises an n -way set associative cache.

16. The processing device of claim 15 wherein said first cache subsystem comprises a direct mapped cache.

17. The processing device of claim 13 and further comprising a plurality of output buffers coupled to outputs of respective data memories.

18. The processing device of claim 17 wherein said output buffers are controlled by said logic.

19. A cache system comprising:

a first n -way cache subsystem, where n is greater than or equal to 1, comprising:

n data memories each having a plurality of entries for storing information from a main memory;
 n tag memories for storing address information identifying a main memory address associated with each of said entries in a corresponding one of said n data memories;

a second m -way cache subsystem, where m is greater than or equal to 1, comprising:

m data memories each having a plurality of entries for storing information from said main memory;
 m tag registers, each storing address information defining a contiguous block of main memory addresses mapped to a corresponding one of said m data memories; and

logic for determining cache hits in said first and second cache subsystems, where hits from said second cache subsystem have precedence over hits from said first subsystem.

20. The cache system of claim 19 and further comprising cache control circuitry for filling one or more lines of one of said cache subsystems after a cache miss in said first and second cache subsystems.

EP 1 107 123 A1

21. The cache system of claim 19 wherein said first cache subsystem comprises an n -way set associative cache.
22. The cache system of claim 21 wherein said first cache subsystem comprises a direct mapped cache.
- 5 23. The cache system of claim 19 and further comprising a plurality of output buffers coupled to outputs of respective data memories.
24. The cache system of claim 23 wherein said output buffers are controlled by said logic.

10

15

20

25

30

35

40

45

50

55




Fig. 1




Fig. 3

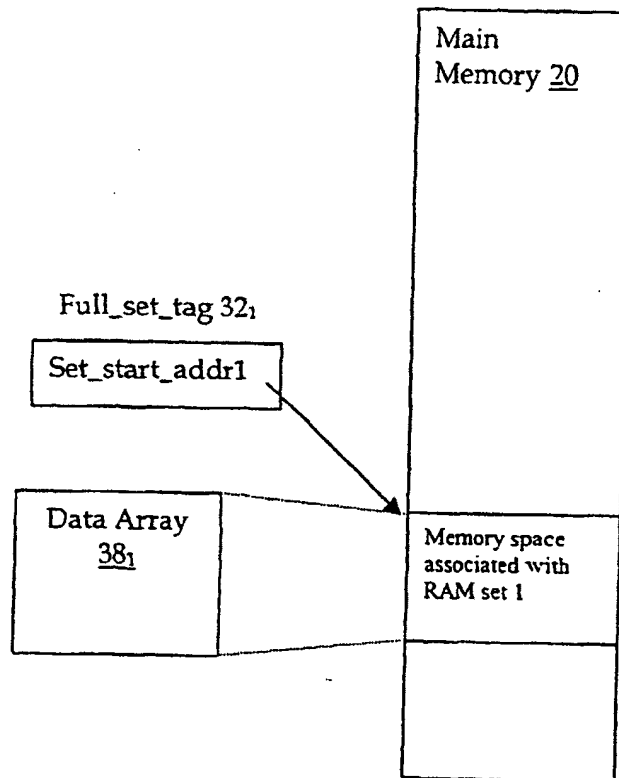
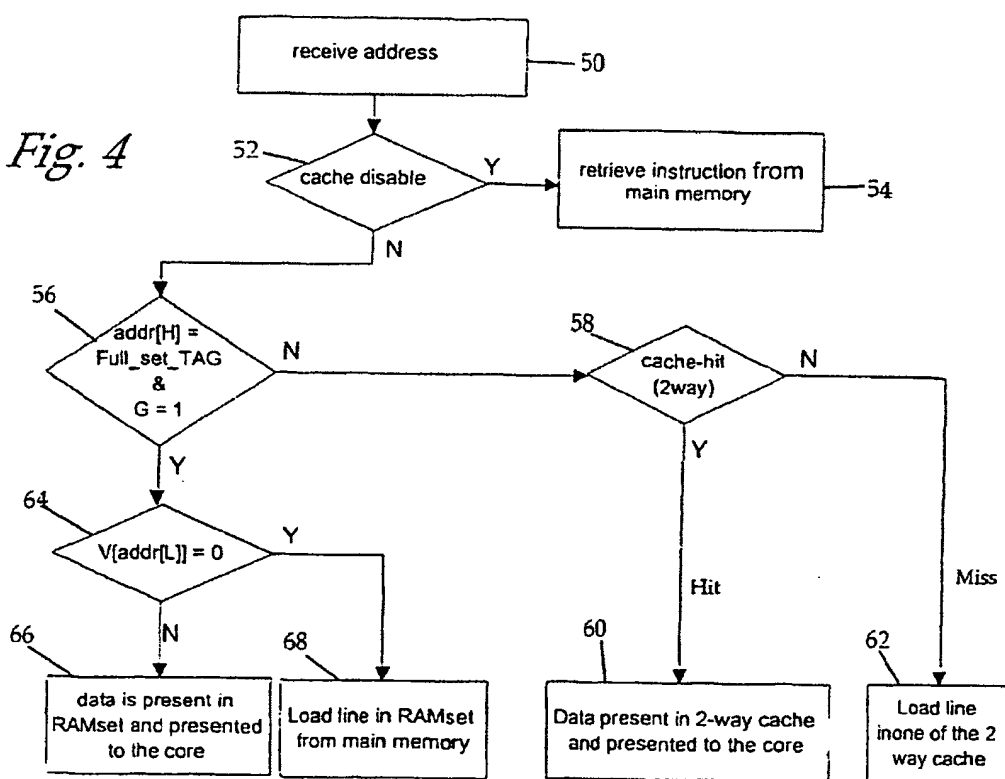


Fig. 4





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 99 40 3042

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	YANG L ET AL: "OPTIMIZING PRIMARY DATA CACHES FOR PARALLEL SCIENTIFIC APPLICATIONS: THE POOL BUFFER APPROACH" , PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON SUPERCOMPUTING, US, NEW YORK, ACM, VOL. CONF. 10, PAGE(S) 141-148 XP000683017 ISBN: 0-89791-803-7	1-3, 6-9, 12-14, 19, 20	G06F12/08
Y	* page 142, left-hand column, paragraph 2 - page 143, left-hand column, line 4; figure 1 *	4, 10, 15, 16, 21, 22	
Y	US 5 014 195 A (FARRELL JAMES A ET AL) 7 May 1991 (1991-05-07) * column 4, line 24 - line 29; figure 1 *	4, 10, 15, 16, 21, 22	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			G06F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 17 April 2000	Examiner Ledrut, P
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EP FORM 1503 03 12 (10/04/01)

EP 1 107 123 A1

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 99 40 3042

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

17-04-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5014195 A	07-05-1991	NONE	

EP FORM P-600

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82